

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Appl. No.	: 09/769,604	Confirmation No.:	4397
Applicant	: Stephen M. Howard, et al.		
Filed	: January 25, 2001		
T.C./A.U.	: 2157		
Examiner	: Ramy M. Osman		
Docket No.	: EMC-002PUS		
Customer No.	: 51576		

AMENDMENT

Mail Stop RCE
Commissioner for Patents
P. O. Box 1450
Alexandria, VA 22313-1450

Dear Sir:

Please amend the above-identified patent application as follows.

Amendments to the Claims are reflected in the listing of claims beginning on page 2 of this paper.

Remarks/Arguments begin on page 6 of this paper.

Amendments to the Claims:

This listing of the claims will replace all prior versions, and listings, of the claims in the application:

1. (Previously Presented) A method of restoring backed up data, comprising:
retrieving, by a data backup storage system, a list of objects that are restorable by a client having a backup/restore module and a logical volume manager to communicate with a storage system, the backup storage system having a storage system interface to communicate with the storage system, a backup storage unit to store backed up data, and a network interface to communicate with the client;
displaying the list of restorable objects for browsing by a user;
generating a list of restorable objects marked for restoration by the user, wherein each of the restorable objects is associated with a particular library;
submitting the list of marked restorable objects for restoration by the client;
executing a restoration of the submitted list of marked restorable objects via a remote procedure call such that multiple restore submissions can be made prior to restore execution.
2. (Original) The method according to claim 1, further including executing multiple restore submissions concurrently.
3. (Original) The method according to claim 1, further including initiating a restore session for the client.
4. (Original) The method according to claim 3, further including creating a restore engine process for the retrieving, browsing, submitting and executing of restore objects.
5. (Original) The method according to claim 4, wherein the client communicates with the restore engine process via remote procedure calls.

6. (Original) The method according to claim 4, wherein the restore engine process is created by a dispatch daemon on a backup storage system server.
7. (Original) The method according to claim 4, wherein the restore engine process is terminated upon completion of the restore execution.
8. (Original) The method according to claim 4, wherein the restore engine process runs on a backup data storage server and further including creating a work item restore process on the backup data server, a server restore process for generating a stream of data to be restored, and a client restore process for receiving the data stream.
9. (Original) The method according to claim 4, further including detecting and identifying libraries that support associated catalogs of backed up data for processing of backed up data by the restore engine process.
10. (Original) The method according to claim 9, further including adding a new library supporting new methods of backing up data.
11. (Original) The method according to claim 9, further including determining object types for backed up data supported by the libraries.
12. (Previously Presented) A method of restoring backed up data, comprising:
 - initiating a restore session for a first client via a dispatch daemon running on a data storage server through a graphical user interface associated with the client;
 - creating a restore engine process in response to a request by the dispatch daemon;
 - establishing a connection between the graphical user interface and the restore engine process;
 - displaying a list of restorable objects for browsing by a user associated with the client via the graphical user interface under the control of the restore engine process;
 - identifying restorable objects marked for restoration by the user under control of the restore engine process;

storing a list of marked restorable objects submitted by the client to the restore engine process; and

executing the restoration of the marked objects under control of the restore engine process independently of the browsing, marking and submitting of the restorable objects such that multiple restore submissions can be made prior to restore execution.

13. (Original) The method according to claim 12, wherein the client communicates with the restore engine process via remote procedure calls.
14. (Original) The method according to claim 12, further including supporting a new backup data method by adding a library corresponding to the new backup data method.
15. (Previously Presented) A data backup and storage system, comprising:
a backup storage system for storing backup data from a client storage system under control of a user associated with the client, the backup storage system including:
a server creating a restore engine process as part of a restore session with a client, the restore engine communicating with the client via remote procedure calls to allow the user to browse restorable objects, mark selected ones of the restorable objects for restoration, submit a list of restorable objects marked by the user, and execute restoration of the submitted list of restorable objects, wherein the restore execution is performed independently of the browse, mark and submit operations such that multiple restore submissions can be made prior to execution of the restore; and
a work item restore process, a server restore process, and a client restore process created by the restore engine process to form a restore triangle for executing the restore operation.
16. (Canceled)
17. (Original) The system according to claim 15, wherein the restore engine process processes libraries upon restore initialization such that libraries can be added to the system for supporting new backup methods.

18. (Original) The system according to claim 17, further including a dispatch daemon for initiating the restore session.
19. (Original) The system according to claim 15, further including further restore engine processes corresponding to further restore sessions initiated by additional clients.
20. (Original) The system according to claim 19, further including additional restore triangles for executing multiple work item restores concurrently.
21. (New) The system according to claim 15, wherein each of the restorable objects is associated with a particular library

REMARKS

Applicant respectfully requests consideration of the above-identified patent application in view of the amendments above and the remarks below.

Claims 1-15 and 17-20 are pending in the application and are rejected. Claim 16 was previously canceled. Claim 21 is herein added.

Upon examining the present application, the undersigned respectfully requests a telephone interview with the Examiner.

The Prior Art Rejections

The Examiner rejects Claims 1-15 and 17-20 under 35 U.S.C. §103(a) over U.S. Patent No. 6,466,952 to Hanes et al. in view of U.S. Patent No. 6,427,149 to Rodriguez.

Applicant submits that the Examiner has construed the claims in an impermissibly and unreasonably broad manner. While it is well established that limitations from the specification are not read into the claims, elements of the claim are to be given their ordinary meaning unless specifically defined by the Applicant. It is also well established that the ordinary meaning of claim terms is determined from the field of art of the invention. Applicant submits that the Examiner has not given claim terms their ordinary meaning and has effectively read out limitations of the claims.

In the present case, as is absolutely clear from the specification and claims, the invention is directed to backup and restore systems, such as system shown in FIG. 3. In such a system, a backup storage system interacts with a client computer and a storage system. This is described in detail in the Background of the Invention, the Summary, and the Detailed Description. Applicant submits that one of ordinary skill in the art of backup and restore systems will not agree that the term backup and storage system can be "interpreted to mean anything that can be used to backup data and act as a storage for data," as asserted by the Examiner.

Similarly, the Examiner asserts that the term library essentially has no meaning in the claim stating that the word library “is broadly interpreted to mean anything that contains a collection of data objects” including a compact disk, or a zip archive file. This is simply not supportable in view of the specification, nor is it the understanding of one of ordinary skill in the art of backup and restore systems. In addition, the Examiner’s broad interpretation is not consistent with the term library in Applicant’s specification. As but one example, page 12 of the specification states that:

“Catalogs contain information associated with backed up data, such as media type, meta data to enabling browsing, bitfile lengths, and rename attributes for a restore. Each catalog is supported by a *library* that interprets the catalog information and passes it to the restore engine process in an expected format. The architecture of the restore system allows new storage architectures to be supported by existing backup storage systems by adding an associated *library* in a predetermined location, such as a library directory. The added *library* provides the catalog information for new types of objects to the restore engine in a usable manner.”

Clearly, the term “library” does not mean “anything that contains a collection of data objects.”

In view of the above, Applicant submits that after giving the claim terms their proper meaning the pending claims are patentably distinguishable over the cited art, as discussed in detail below.

Claim 1 requires a method of restoring backed up data, including retrieving, by a data backup and storage system, a list of objects that are restorable by a client, displaying the list of restorable objects for browsing by a user, generating a list of restorable *objects marked for restoration* by the user, wherein *each of the restorable objects is associated with a particular library*, submitting the list of marked restorable objects for restoration by the client, and executing a restoration of the submitted list of marked restorable objects via a remote procedure call such that *multiple restore submissions can be made prior to restore execution*.

Hanes discloses transferring data from old media to new media. For example, Hanes teaches migrating data from a floppy disc to a CD-RW disc. While Hanes teaches that the user can select source files for migration to the destination media, Hanes does not teach or suggest that restorable objects are associated with a particular library, as claimed. As described in Applicant’s specification at page 12, first full paragraph,

“In general, libraries to support additional catalogs associated with new backup methods and new data storage types can be added with minimal overall impact on the restore system. Catalogs contain information associated with backed up data, such as media type, meta data to enabling browsing, bitfile lengths, and rename attributes for a restore. Each catalog is supported by a library that interprets the catalog information and passes it to the restore engine process in an expected format. The architecture of the restore system allows new storage architectures to be supported by existing backup storage systems by adding an associated library in a predetermined location, such as a library directory. The added library provides the catalog information for new types of objects to the restore engine in a usable manner.”

Applicant submits that Hanes does not teach or suggest associating restorable objects with a particular library, as claimed.

In addition, Hanes appears limited to a personal computer while claim 1 requires a data backup and storage system. Applicant submits that Hanes is not a backup and storage system as understood in the art. In the Background of the Invention section of Applicant's specification, a Symmetrix storage system is listed as an exemplary storage system that can be connected to a client. Hanes simply cannot perform as a data backup and storage system, such as the EMC Symmetric system.

Applicant submits that Rodriguez fails to overcome any of the deficiencies of Hanes set forth above.

Rodriguez merely teaches a technique to enable an Internet user to utilize a browser application to see and select files in a ZIP archive file by providing hyper text links in an HTML document. Rodriguez avoids the need for a user to download the entire ZIP archive file. Applicant submits that the trivial archive file recovery technique taught by Rodriguez is quite irrelevant to backing up huge amounts of data contained in submit objects, which can include databases, and restoring the backed up data using a data backup system, such as the one shown and described in Applicant's specification.

Applicant submits that Rodriguez fails to teach or suggest any of the claimed method of restoring backed up data, which requires, among other things, generating a list of restorable

objects marked for restoration by the user, wherein *each of the restorable objects is associated with a particular library*, and executing a restoration of the submitted list of marked restorable objects via a remote procedure call such that *multiple restore submissions can be made prior to restore execution*.

Accordingly, Applicant submits that claim 1 is patentably distinguishable over Hanes and/or Rodriquez. For at least substantially the same reasons, Applicant submits that claims 2-15 and 17-20 are also distinguishable over the cited references.

Notwithstanding the above, Applicant submits that certain dependent claims are patentably distinguishable over the cited art for additional reasons.

Claim 2, for example, requires, “executing multiple restore submissions concurrently.”

The Examiner points to col. 5, lines 1-15:

“Preferably, data transfer application 20 allows the user to select groups of files to be transferred to different directories on the second storage media 18. In other words, the user could select a first group of files contained on the source storage media 14 to be transferred to one directory on the destination storage media 18, a second group of files contained on the source storage media 14 to be transferred to a different directory (or subdirectory) on the destination storage media 18, and so on. This capability enables disc organization by allowing the user to group files located on one disc and automatically move them to different locations on another disc.

When the files are selected to be transferred, a subdirectory containing the name of the identifier is created on the second destination media 18 the selected files are then transferred 26 to the destination media 18 under the newly created subdirectory identified by the identifier.”

The Examiner also points to col. 6, lines 20-35:

“Alternatively, the data transfer function 35 itself implements this functionality.

Preferably, file selection function 32 provides the ability to select different groups of files to be transferred to different directories on the second storage media 18. In this case, the user could select a first group of files contained on the first storage media 14 to be transferred to one directory (identified by a first user selected identifier) on the second storage media 18, a second group of files contained on the first storage media 14 to be transferred to a different directory (identified by a second user selected identifier) on the second storage media 18, and so on. This preferred embodiment of file selection function 32 enables disc organization by

allowing the user to group files on a disc and automatically move them to different locations on the other disc.”

While the above passages may disclose allowing a user to select multiple files, Applicant submits that there is no teaching or suggest of *concurrent execution* of multiple restores, as claimed.

Claim 9, for example, requires “*detecting and identifying libraries* that support associated *catalogs of backed up data* for processing of backed up data by the restore engine process.” The Examiner points to col. 6, lines 8-25 and 49-67, which is set forth below:

“File selection function 32 allows the user to select a group of files to transfer to the second storage media 18 present in the destination storage device 16. If source storage device 12 is a tape drive, data transfer function 35 may invoke commercial tape backup/restoration software to mount the tape, extract volume information, position the tape to the desired volume, extract file information from the desired volume, and return the file information to the file selection function 32 for presentation to the user. Once the user selects the desired files from the presented available files, data transfer function 35 invokes the restore function of the commercial tape backup/restoration software to restore the selected files to local memory (i.e., the hard drive 10). Alternatively, the data transfer function 35 itself implements this functionality.

Preferably, file selection function 32 provides the ability to select different groups of files to be transferred to different directories on the second storage media 18. In this case, the user could select a first group of files contained on the first storage media 14 to be transferred to one directory....

When the files are selected to be transferred, data transfer application 20 preferably indexes 25 the contents of the selected files. This is performed by indexing function 33. Preferably, when the files to be transferred are selected, an index 3 describing the contents of the selected files is created by the indexing function 33. For example, if the selected files to be transferred are word processing files, the contents of the index 3 would include a set of keys comprising information about the contents of the word processing files. The index 3 is preferably stored along with the selected files on the destination media 18. Then, at a later time, if the user wants to search for files containing certain text, data transfer application 20 searches the contents of the index 3 to find all the files that have corresponding keys that match one or more of the searched for text. In the preferred embodiment, the index 3 is also stored on the local hard drive 10 of the computer system 2. Storing the index 3 on the local hard drive 10 allows the user to perform a quick search for the location of and contents of a file without actually mounting...”

In reviewing the above passage, Applicant respectfully submits that Hanes does not teach or suggest the subject matter of claim 9.

Claim 10 requires “adding a new library supporting new methods of backing up data” and claim 11 requires “determining object types for backed up data supported by the libraries” for which the Examiner points to the same passage as for claim 9. Applicant respectfully submits that Hanes does not teach or suggest the claimed subject matter.

In view of the above, Applicant submits that claims 1-15 and 17-20 are patentably distinguishable over the cited references.

The Examiner is respectfully invited to telephone the undersigning attorney if there are any questions regarding this Amendment or this application.

Applicant does not acquiesce to any assertion made by the Examiner that is not specifically addressed herein.

The Assistant Commissioner is hereby authorized to charge payment of any additional fees associated with this communication or credit any overpayment to Deposit Account No. 500845.

Dated: 27 Apr 06

Respectfully submitted,

DALY, CROWLEY, MOFFORD & DURKEE, LLP

By: 

Paul D. Durkee

Reg. No. 41,003

Attorney for Applicant(s)

354A Turnpike Street - Suite 301A

Canton, MA 02021-2714

Tel.: (781) 401-9988,

Fax: (781) 401-9966

pdd@dc-m.com

28064